arXiv:1608.04728v1 [cs.IT] 16 Aug 2016

A Practical Study of Longitudinal Reference Based Compressed Sensing for MRI

Samuel Birns, Bohyun Kim, Stephanie Ku, Kevin Stangl, Deanna Needell

August 17, 2016

Abstract

Compressed sensing (CS) is a new signal acquisition paradigm that enables the reconstruction of signals and images from a low number of samples. A particularly exciting application of CS is Magnetic Resonance Imaging (MRI), where CS significantly speeds up scan time by requiring far fewer measurements than standard MRI techniques. Such a reduction in sampling time leads to less power consumption, less need for patient sedation, and more accurate images. This accuracy increase is especially pronounced in pediatric MRI where patients have trouble being still for long scan periods. Although such gains are already significant, even further improvements can be made by utilizing past MRI scans of the same patient. Many patients require repeated scans over a period of time in order to track illnesses and the prior scans can be used as *references* for the current image. This allows samples to be taken *adaptively*, based on both the prior scan and the current measurements. Work by Weizman [20] has shown that so-called *reference based adaptive-weighted temporal Compressed Sensing MRI* (LACS-MRI) requires far fewer samples than standard Compressed Sensing (CS) to achieve the same reconstruction signal-to-noise ratio (RSNR). The method uses a mixture of reference-based and adaptive-sampling. In this work, we test this methodology by using various adaptive sensing schemes, reconstruction methods, and image types. We create a thorough catalog of reconstruction behavior and success rates that is interesting from a mathematical point of view and is useful for practitioners. We also solve a grayscale compensation toy problem that supports the insensitivity of LACS-MRI to changes in MRI acquisition parameters and thus showcases the reliability of LACS-MRI in possible clinical situations.

1 Introduction

Magnetic Resonance Imaging (MRI) has become a staple of modern medicine. Because it is noninvasive and it produces high quality images, MRI is used widely in both treatment and research [9]. However, a typical MRI can take upwards of an hour and requires the patient to be completely still; even the slightest movement will induce enough noise in the image to make the reconstruction unusable [11]. For this reason, advancements that decrease the amount of time a MRI image takes are incredibly beneficial. MRI is an ideal application of Compressed Sensing (CS), a method that enables signal reconstruction with only a fraction of the original information by numerically solving under-determined linear systems [3]. Since we only need to collect a small percentage of the original data, CS can make MRI much faster [8]. This is possible because MRI images can be represented in a *sparse* way; that is, there are linear maps that transform these images into matrices with only a small number of non-zero entries.

1.1 Mathematical Formulation

CS relies on solving on under-determined linear systems using mathematical optimization. Consider a signal (e.g. vector, matrix or image) $\boldsymbol{x} \in \mathbb{C}^n$ that has $s \ll n$ non-zero entries. We call such a signal s-sparse, and write $\|\boldsymbol{x}\|_0 := |\{i : \boldsymbol{x}_i \neq 0\}| = s << n$. We take (noisy) linear measurements of $\boldsymbol{x}$ of the form $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}$, where $\boldsymbol{A}$ is an $m \times n$ measurement matrix, $\boldsymbol{e} \in \mathbb{C}^m$ is a gaussian noise term, and $\boldsymbol{y} \in \mathbb{C}^m$ is called the measurement vector. When $m \ll n$, $\boldsymbol{y}$ can be used to efficiently recover most of the information in $\boldsymbol{x}$. The goal of CS is to (i) design efficient measurement matrix types $\boldsymbol{A}$ that capture the most important information in desired signals, and (ii) design reconstruction methods that recover the signal $\boldsymbol{x}$ from its measurements $\boldsymbol{y}$. Since the recovery problem is highly

underdetermined, identifying the true signal $\boldsymbol{x}$ is non-trivial, and only becomes a well-posed problem under the assumption of sparsity. For example, whenever $\boldsymbol{A}$ is one-to-one on sparse signals and there is no noise in the measurements (i.e. $\boldsymbol{e} = 0$), one can solve the ℓ_0-minimization problem to accurately recover the signal $\boldsymbol{x}$:

$$\min_{\boldsymbol{g}\in\mathbb{C}^N} \left\|\boldsymbol{g}\right\|_0 \text{ s.t. } \boldsymbol{y} = \boldsymbol{A}\boldsymbol{g}. \tag{1}$$

Unfortunately in general, this problem is NP-Hard [15], so more practically useful recovery methods are needed.

Candès et. al. [4] have developed two fundamental results that link sparsity to theoretical image reconstruction. The authors introduced the Restricted Isometry Property (RIP) [1]. A matrix is said to satisfy the RIP if it is nearly isometric (i.e. distance preserving) on sparse vectors; specifically, an $m \times n$ matrix $\boldsymbol{A}$ satisfies the RIP (of order s) if there exists $\delta_s \in (0,1)$ such that for all $m \times s$ submatrices $\boldsymbol{A}_s$ of $\boldsymbol{A}$ and for every s-sparse vector $\boldsymbol{x}$,

$$(1-\delta_s)\|\boldsymbol{x}\|_2^2 \leq \|\boldsymbol{A}_s\boldsymbol{x}\|_2^2 \leq (1+\delta_s)\|\boldsymbol{x}\|_2^2 \tag{2}$$

An early result in CS [2, 18] states that, given a set of randomly selected indices $\boldsymbol{\Omega}$ from $\{1,2,\ldots,n\}$ with $|\boldsymbol{\Omega}| \approx s\log^4 N$, if $\boldsymbol{F}_u$ is the subsampled Fourier transform (i.e. the $m \times n$ submatrix of the discrete Fourier transform (DFT) whose rows correspond to the indices in $\boldsymbol{\Omega}$), then with high probability $\boldsymbol{F}_u$ satisfies the RIP, and the optimization problems

$$\min_{\boldsymbol{g}\in\mathbb{C}^N} \left\|\boldsymbol{g}\right\|_0 \text{s.t. } \boldsymbol{g}|_{\boldsymbol{\Omega}} = \boldsymbol{F}_u \quad \text{and} \quad \min_{\boldsymbol{g}\in\mathbb{C}^N} \left\|\boldsymbol{g}\right\|_1 \text{s.t. } \boldsymbol{g}|_{\boldsymbol{\Omega}} = \boldsymbol{F}_u \tag{3}$$

are equivalent in the sense that these optimization problems yield the same solutions. Since ℓ_1 optimization can be cast as an efficient linear programming problem, the ability to solve ℓ_1 optimization in lieu of ℓ_0 optimization has made CS much more practically useful. In addition, Candès et. al. [5] show that if a matrix A satisfies the RIP and if a vector $\boldsymbol{x}_0 \in \mathbb{R}^m$ is s-sparse, given noisy measurements $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{e}$, the solution $\hat{\boldsymbol{x}}$ to the optimization problem

$$\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \text{ s.t. } \left\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{y}\right\|_2 \leq \epsilon, \tag{4}$$

satisfies

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}_0\|_2 \leq C \cdot \epsilon, \tag{5}$$

where $\epsilon \geq \|e\|_2$. The combination of the two previous equations gives CS its utility since ℓ_1-minimization produces accurate signal recovery from highly underdetermined measurements.

Closely related to the RIP is the *coherence* of a matrix. Intuitively, the coherence of a matrix is a measure of how similar its columns are; the higher the coherence, the more similar its columns. Formally, coherence of a matrix $\boldsymbol{A}$ is defined as

$$\mu(\boldsymbol{A}) = \max_{j<k} \frac{|\langle \boldsymbol{A}_j, \boldsymbol{A}_k\rangle|}{\left\|\boldsymbol{A}_j\right\|_2 \|\boldsymbol{A}_k\|_2} \tag{6}$$

where $\langle\cdot\rangle$ denotes the standard inner product, and we write $\boldsymbol{A}_j$ to denote the jth column of $\boldsymbol{A}$. Note that $0 \leq \mu(\boldsymbol{A}) \leq 1$ for any matrix $\boldsymbol{A}$ and that the closer $\mu(\boldsymbol{A})$ is to 0 (or the more *incoherent* $\boldsymbol{A}$ is), the closer the columns of $\boldsymbol{A}$ are to being orthogonal. Thus, coherence - or, more accurately, incoherence - and the RIP are very closely related.

1.2 Sparse representations

The above theory and formulation hold for signals which are sparse. However, most signals in practice are not themselves sparse, but rather can be represented sparsely in some basis. For example, natural images are sparse in wavelet bases [6]; that is, we can represent an image $\boldsymbol{x}$ by $\boldsymbol{x} = \boldsymbol{W}\boldsymbol{z}$ for some sparse (or nearly sparse) signal $\boldsymbol{z}$ and an orthonormal matrix $\boldsymbol{W}$. Since for any Gaussian matrix $\boldsymbol{A}$ and orthonormal basis $\boldsymbol{W}$, the product $\boldsymbol{A}\boldsymbol{W}$ is

also Gaussian, the results above hold for Gaussian measurement matrices even when the signal is represented sparsely in an orthonormal basis. For other types of RIP matrices, one requires that the RIP measurement matrix and the sparsifying orthonormal basis be incoherent [17]. In the MRI setting, the measurement matrix is a subsampled Fourier matrix and the sparsifying basis is typically a wavelet basis; unfortunately, these ensembles are not incoherent. To overcome this problem and prove theoretical bounds on reconstruction while using wavelet bases, Krahmer and Ward [12] use an approach known as *variable density sampling* (VDS), which we describe in detail below. This type of sampling stems from the idea that low frequency components contain the majority of the information about the underlying signal, so sampling primarily near the origin in k-space[1] is more efficient and accurate than uniform random subsampling [13]. Specifically, Krahmer and Ward [12] define a novel probability density function that relates frequencies and k-space coordinates and is maximized at the origin. They then showed that the error from the reconstruction using this pdf is bounded.

1.3 Longitudinal MRI

While there is currently a wide variety of literature on CS and its uses in MRI, most of it does not take into account that often a patient who undergoes an MRI will undergo several MRIs spaced out over regular intervals of time. Since these time intervals are relatively short, i.e. often less than a month, it can be expected that there is little baseline change from one scan to the next. Weizman et. al. exploit this by introducing a reference image (a previous MRI scan) to classical CS reconstruction techniques [20]. Specifically, Weizman et. al. iteratively solve:

$$\min_{\boldsymbol{x}} \|\boldsymbol{W}_1\boldsymbol{\Psi}\boldsymbol{x}\|_1 + \lambda\|\boldsymbol{W}_2(\boldsymbol{x}-\boldsymbol{x}_0)\|_1 \quad \text{s.t.} \quad \|\boldsymbol{F}_u\boldsymbol{x}-\boldsymbol{y}\|_2 < \epsilon, \tag{7}$$

where $\boldsymbol{x}_0$ is a reference image, $\boldsymbol{y} = \boldsymbol{F}_u\boldsymbol{x}_\star$ is the (possibly noisy) measurement vector for the current image $\boldsymbol{x}_\star$, $\boldsymbol{\Psi}$ is a sparsifying operator such as the wavelet transform, $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$ are diagonal weighted matrices (to be described momentarily), $\boldsymbol{F}_u$ is the undersampled Fourier transform operator, and λ is a parameter that can be modified to either emphasize sparsity (by making $\|\boldsymbol{\Psi}\boldsymbol{x}\|_1$ dominant in the minimization problem) or fidelity with the ground truth (by making $\|\boldsymbol{x}-\boldsymbol{x}_0\|_1$ dominant). To solve the minimization problem, Weizman et. al. define an algorithm which they refer to as Longitudinal Adaptive Compressed Sensing MRI (LACS-MRI).

In each step of LACS-MRI, (7) is solved and new samples are selected using the discrete probability density function

$$f_S(k_y) = \gamma f_{ND}(k_y) + (1-\gamma) f_{VD}(k_y). \tag{8}$$

f_{VD} is a variable density sampling probability density function that is maximized at $k_y = 0$. f_{ND} is a density function designed to locate missing frequencies in the current signal estimation and $0 \le \gamma \le 1$ is a fidelity term that relates the reference and current image. The weighting matrices $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$ are updated based on the reconstruction in order to emphasize important signal entries, and γ is updated based on the entries of $\boldsymbol{W}_2$.

By construction, when $\gamma \approx 0$, the f_{VD} term dominates and a new set of k-space coordinates near the origin is sampled; conversely, when $\gamma \approx 1$, the f_{ND} term dominates and the next iteration of k-space coordinates are taken from almost exactly the same region as the previous set of k-space coordinates. LACS-MRI updates γ such that $\gamma \approx 0$ if the reconstruction is far from the ground truth and $\gamma \approx 1$ if the reconstruction is close to the ground truth.

LACS-MRI is innovative since it uses both variable density sampling and adaptive sampling, where adaptive sampling is simply the process of choosing samples based off previous reconstructions, as opposed to sampling the same values regardless of previous measurements or iterations. While adaptive sampling may appear to be a logical way to drastically improve reconstruction, Davenport et. al. [7] shows that, perhaps counterintuitively, the performance of adaptive sampling is theoretically limited. In fact, they show that even in the best case, adaptive sampling reconstruction has an error term roughly of $s\log^2 n$, where s is the sparsity of the target image and n is the dimension of the target image, which is far lower than was previously conjectured. However, Davenport

[1] We use the k-space representation of frequency space, where the low frequencies reside at the origin. Subsampling a discrete Fourier ensemble is then equivalent to sampling coordinates k_y from this k-space.

et. al. also show that, despite this theoretical result, when samples are forced to be chosen from a set of predetermined measurements, a property they refer to as *constrained* measurements, adaptive sampling does indeed offer significant improvement over nonadaptive sampling.

Given an estimate of the signal's support, they provide an optimization problem that adaptively selects the optimal measurements from the Fourier ensemble (i.e. the optimal k-space frequencies to be used for the measurements). This work is particularly relevant here as MRI measurements are inherently constrained, since measurements must be taken in k-space. In this paper, we explore this type of adaptive sampling within the LACS-MRI framework, and compare various sampling techniques and reconstruction approaches.

1.4 Contributions

We expand upon the work above by investigating the effects of combining variable density sampling and adaptive sampling into one algorithm. We study this algorithm empirically by testing its performance under various sampling schemes. Additionally, we implement several modifications motivated by practical applications, including an approach that addresses different grayscale calibrations across MRI scans.

2 LACS-MRI

In this section, necessary background and details will be developed about LACS-MRI. LACS-MRI uses past reference images to help reconstruct the follow-up images for patients who require multiple brain scans. If the previous and current scans are similar, we may accelerate the recovery of follow-up scan and improve the accuracy [19]. Sometimes the reference image is not similar to the follow-up image due to surgery or treatment. In most cases the underlying structure of the brain does not change very much so we can still benefit from the reference image. Recall that the algorithm provides a reconstruction by solving the convex optimization problem:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \underbrace{\|\boldsymbol{W}_1\boldsymbol{\Psi}\boldsymbol{x}\|_1}_{\text{term 1}} + \underbrace{\lambda\|\boldsymbol{W}_2(\boldsymbol{x}-\boldsymbol{x}_0)\|_1}_{\text{term 2}} \text{ s.t. } \|\boldsymbol{F}_u\boldsymbol{x}-\boldsymbol{y}\|_2 < \epsilon, \tag{9}$$

where $\boldsymbol{x}_0$ is the reference image, $\hat{\boldsymbol{x}}$ is our estimation of the follow-up brain image, and $\boldsymbol{y}$ is the vector of complex measurements in k-space of the true follow-up brain image $\boldsymbol{x}_\star$. $\boldsymbol{\Psi}$ is a sparsifying wavelet transform matrix and the $\boldsymbol{F_u}$ term is the sampling operator, which in our case represents line samples of k-space in 2D imaging. $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$ are diagonal weight matrices and their ith entries, $[\cdot]_i$, are calculated as follows:

$$w_1^i = \begin{cases} 1, & \text{if } \frac{[|\boldsymbol{\Psi}(\hat{\mathbf{x}}-\mathbf{x}_0)|]_i}{1+[|\boldsymbol{\Psi}(\hat{\mathbf{x}}-\mathbf{x}_0)|]_i} > \epsilon_1 \\ \frac{1}{1+[|\boldsymbol{\Psi}\mathbf{x}_0|]_i}, & \text{otherwise} \end{cases} \tag{10}$$

$$w_2^i = \frac{1}{1+[|\hat{\boldsymbol{x}}-\boldsymbol{x}_0|]_i} \tag{11}$$

where $\epsilon_1 > 0$ is some fixed threshold parameter.

In (9), term 1 enforces the sparsity constraint, while term 2 enforces the similarity of the $\hat{\boldsymbol{x}}$ to the reference scan $\boldsymbol{x}_0$. Diagonal matrices $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$ are iterative weighting terms that depend on the difference between the current reconstruction and the reference image. Observe that if the difference between the ith component of $\hat{\boldsymbol{x}}$ and $\boldsymbol{x}_0$ is large, $w_1^i \approx 1$ and $w_2^i \approx 0$. In other words, when the difference between the reference image and the current scan is large, the optimization problem (9) becomes conventional ℓ_1 minimization with wavelet sparsity (L1-W):

$$\min_{\boldsymbol{x}} \|\boldsymbol{\Psi}\boldsymbol{x}\|_1 \text{ s.t. } \|\boldsymbol{F}_u\boldsymbol{x}-\boldsymbol{y}\|_2 < \epsilon. \tag{12}$$

On the other hand, when $\hat{\boldsymbol{x}}$ and $\boldsymbol{x}_0$ are reasonably close, the ith entry of $\boldsymbol{W}_1$ takes a large value to enforce sparsity if the ith entry of $\boldsymbol{\Psi}\boldsymbol{x}_0$ is small, or relaxes the sparsity condition if the ith entry of $\boldsymbol{\Psi}\boldsymbol{x}_0$ is large. $\boldsymbol{W}_2$ behaves in a similar way. It enforces the similarity between $\hat{\boldsymbol{x}}$ and $\boldsymbol{x}_0$ for the regions that are highly similar while

it relaxes the constraint for the regions that are not similar. λ is a regularization parameter that controls the relative importance of term 1 and term 2.

In LACS-MRI, we select measurements using mixed PDFs of variable density and adaptive density functions. In particular, we will utilize the following formal definitions of sampling densities:

$$f_{VD}(k_y) = \frac{(1-\frac{2}{n}|k_Y|)^p}{\sum_{k_y}(1-\frac{2}{n}|k_y|)^p} \tag{13}$$

$$f_{ND}(k_y) = \frac{g_{ND}(k_y)}{\sum_{k_y} g_{ND}(k_y)}, \qquad g_{ND}(k_y) = \frac{\sum_{i\in k_y}[|\boldsymbol{F}\hat{\boldsymbol{x}} - \boldsymbol{F}\boldsymbol{x}|]_i}{[|\boldsymbol{F}\hat{\boldsymbol{x}}| + |\boldsymbol{F}\boldsymbol{x}|]_i} \tag{14}$$

$$f_R(k_y) = \frac{g_R(k_y)}{\sum_{k_y} g_R(k_y)}, \qquad g_R(k_y) = \sum_{i\in k_y}[|\boldsymbol{F}\boldsymbol{x}_0|]_i. \tag{15}$$

We may then mix these PDFs as described in (8). There, γ is a fidelity term that controls the weight of variable and adaptive density sampling, where $\gamma = \frac{1}{N}\sum_{i=1}^{n} w_2^i$. This term increases when $\hat{\boldsymbol{x}}$ and $\boldsymbol{x}$ are similar, which leads to a heavier weighting of adaptive density sampling. Intuitively, it makes sense that we want to use the adaptive density sampling more than the variable density sampling when the reference image is close to the follow-up image since we expect most of the signal information to reside within reference image.

We will also consider other sampling strategies proposed in various works. Krahmer et. al. suggest a nearly optimal Variable Density Sampling (f_{VDS}) function that is able to reconstruct images using various sparsification methods [12]. f_{VDS} is a variable density sampling function that selects 2D k-space coordinate pairs (k_1, k_2) with probability:

$$\text{Prob}[(\omega_1^j, \omega_2^j) = (k_1, k_2)] = C' \min\left(C, \frac{1}{(k_1^2 + k_2^2)^p}\right), \tag{16}$$

where p is a user selected parameter (typically $0 < p < 2$), while C and C' denote appropriate constants. Davenport et. al. suggest the optimal Constrained Adaptive Sensing (f_A) function to select measurements based on a signal estimation [7]. f_A is defined by taking a sequence of m rows from $\boldsymbol{F}$, where $\boldsymbol{F}$ is the discrete Fourier transform matrix (DFT) that transforms a vectorized image [7]. Letting $\boldsymbol{a}_i$ be the ith row of the discrete Fourier transform $\boldsymbol{F}$, a sequence of $\boldsymbol{a}_i$ can be selected by the following (where $\boldsymbol{A}'$ denotes the submatrix with rows $\boldsymbol{a}_i$):

$$\{\boldsymbol{a}_i\}_{i=1}^m = \underset{\{\{\boldsymbol{a}_i\}_{i=1}^m\}}{\arg\min}\ \text{tr}(((\boldsymbol{A}'\boldsymbol{\Psi}_\Lambda) * \boldsymbol{A}'\boldsymbol{\Psi}_\Lambda)^{-1}) \tag{17}$$

However, this is intractable, so Davenport et. al. [7] consider the relaxation:

$$\hat{\boldsymbol{S}} = \underset{\text{diagonal matrices } \boldsymbol{S}\geq 0}{\arg\min}\ \text{tr}(((\boldsymbol{F}\boldsymbol{\Psi}_\Lambda) * \boldsymbol{S}\boldsymbol{F}\boldsymbol{\Psi}_\Lambda)^{-1}) \quad \text{subject to } \text{tr}(\boldsymbol{S}) \leq \varepsilon, \tag{18}$$

where $\hat{\boldsymbol{S}}$ is a diagonal matrix whose diagonal entries correspond to how much of the corresponding Fourier measurement should be utilized. Then, the probability of selecting the ith row of $\boldsymbol{F}$ can be calculated by:

$$f_A(i) = \frac{\hat{S}_{ii}}{m}. \tag{19}$$

The pseudocode for LACS-MRI as presented in [19] is given below. We next turn to our numerical investigations of this method combined with various sampling techniques.

Algorithm 1 LACS-MRI

Input: $\boldsymbol{x}_0$ (the reference image), $\boldsymbol{N}$ (the number of iterations), η (total percentage k-space to be sampled)
Output: $\hat{\boldsymbol{x}}$ (estimated image)
Let $\boldsymbol{W}_1 = I, \boldsymbol{W}_2 = 0$, and $\boldsymbol{y} = 0$. Pick $k = \frac{\eta}{N}$ sampling locations randomly (using VDS or uniformly).
Sampling: for i=1 to $\boldsymbol{N}$
  Sample k locations, $\boldsymbol{y}_i = (\boldsymbol{F}_u)_i \boldsymbol{X}$
  Let $\boldsymbol{y} = \boldsymbol{y} + \boldsymbol{y}_i$
  Solve $\hat{\boldsymbol{x}} = \min_{\boldsymbol{x}} \|\boldsymbol{W}_1 \boldsymbol{\Psi} \boldsymbol{x}\|_1 + \lambda \|\boldsymbol{W}_2 (\boldsymbol{x} - \boldsymbol{x}_0)\|_1 + \|\boldsymbol{F}_u \boldsymbol{x} - \boldsymbol{y}\|_2$
  Update $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$ via (10) and (11). Use the sampling PDF (8) to select k new sampling locations.

3 Experimental Results

3.1 Reconstruction performance with various PDFs

We compare the image recovery of ℓ_1 minimization with wavelet sparsity (L1-W) and Adaptive Weighted Temporal CS-MRI (LACS-MRI) algorithms through experiments. L1-W is a conventional compressed sensing problem given in (12). We solve this problem using the Projection Over Convex Sets (POCS) method, which was shown to be faster and to have a better recovery in preliminary experiments [14]. Weizman et.al. have tested the LACS-MRI algorithm with Variable Density Sampling for LACS-MRI (f_{VD}), Reference LACS-MRI sampling (f_{ND}), and Adaptive Density Sampling for Reference (f_R) [19]. However, these sampling functions have not been proven to be optimal. These algorithms can be improved by utilizing other variable density or adaptive density sampling functions.

We consider twenty two combinations of experiments in Table 1 and measure their mean Reconstruction Signal to Noise Ratios (RSNR) over 50 iterations to decide the best scheme. The mixture of variable density sampling and adaptive density sampling is defined by (8) with a fidelity term γ to select measurements in k-space in the LACS-MRI method. The idea of creating a mixed PDF can be applied to classical L1-minimization as well. Thus, we experimented with both LACS-MRI and classical L1-minimization (L1-W) along with two variable density sampling PDFs, f_{VD} and f_{VDS}, and three different adaptive density sampling PDFs, f_R, f_A, and f_{ND} to create mixed PDFs as in (8). These PDFs are described in (13), (16), (15), (19), and (14), respectively. For some of the tests, we purely used adaptive density sampling or variable density sampling. For each unique PDF defined, we experimented with a 32 x 32 phantom and a 512 x 512 brain image according to Table 1. For each true image $\boldsymbol{x}_\star$ and recovered estimation $\hat{\boldsymbol{x}}$, we calculated the Reconstruction SNR defined by RSNR = $20\log(\|\boldsymbol{x}_\star\|_2 / \|\boldsymbol{x}_\star - \hat{\boldsymbol{x}}\|_2)$, over 50 trials and calculated the mean over all 50 trials (MSNR). We also varied the compression level $\eta = \frac{m}{n}$, where m is the number of measurements, and n is the dimension of the image. We tested each of the cases, varying η from 0.03-0.21 in increments of 0.03 to observe the dynamics of the MSNR.

Table 1: Case Reference Table

Cases	*Variable Density Sampling*	*Adaptive Density Sampling*	*Algorithm*
case 1	f_{VD}	f_R	LACS-MRI
case 2	f_{VDS}	f_R	LACS-MRI
case 3	f_{VD}	-	LACS-MRI
case 4	f_{VDS}	-	LACS-MRI
case 5	-	f_R	LACS-MRI
case 6	-	f_A	LACS-MRI
case 7	f_{VD}	f_A	LACS-MRI
case 8	f_{VDS}	f_A	LACS-MRI
case 9	f_{VD}	f_R	L1-W
case 10	f_{VDS}	f_R	L1-W
case 11	f_{VD}	-	L1-W
case 12	f_{VDS}	-	L1-W
case 13	-	f_R	L1-W
case 14	-	f_A	L1-W
case 15	f_{VD}	f_A	L1-W
case 16	f_{VDS}	f_A	L1-W
case 17	f_{VD}	f_{ND}	LACS-MRI
case 18	f_{VDS}	f_{ND}	LACS-MRI
case 19	-	f_{ND}	LACS-MRI
case 20	f_{VD}	f_{ND}	L1-W
case 21	f_{VDS}	f_{ND}	L1-W
case 22	-	f_{ND}	L1-W

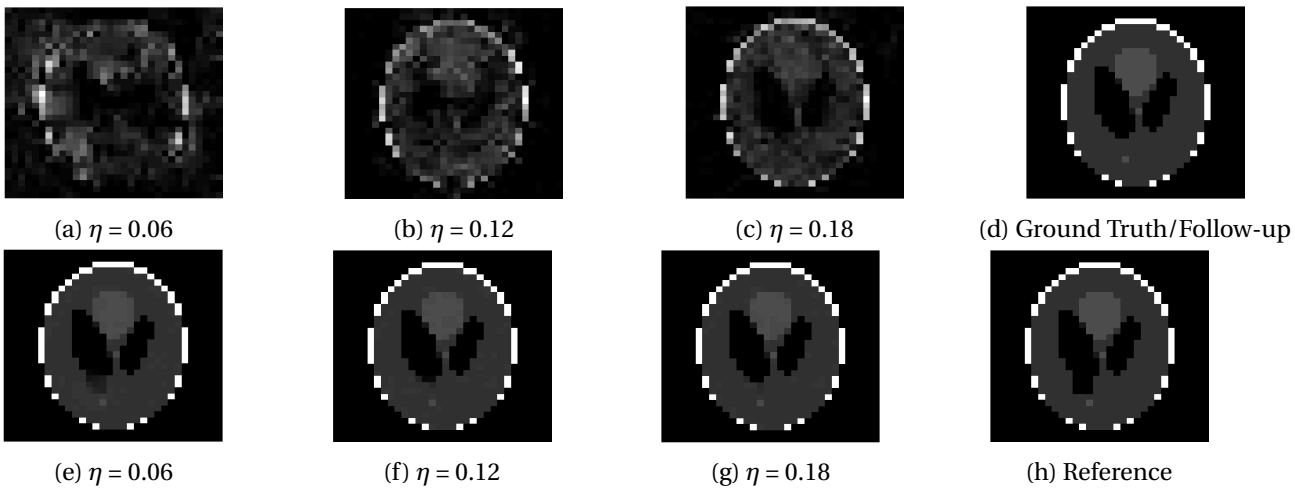

(a) $\eta = 0.06$ (b) $\eta = 0.12$ (c) $\eta = 0.18$ (d) Ground Truth/Follow-up

(e) $\eta = 0.06$ (f) $\eta = 0.12$ (g) $\eta = 0.18$ (h) Reference

Figure 1: Image reconstruction from case 13 (a, b, and c) and case 1 (e, f, and g) of a phantom image. Case 13 is the best performing algorithm that uses L1-W and case 1 is the best performing algorithm that uses LACS-MRI

Figure 3 displays the reconstruction MSNR for the various test cases as cataloged in Table 1. Figures 1 and 2 display some examples of reconstructed images for cases of interest.

The image recovery results for the phantom image experiment were quite extreme, as shown in Figure 3. We split all the cases into two groups: a group with high MSNR and a group with low MSNR. While grouping these cases we noticed that all of the cases that used LACS-MRI were classified as high MSNR, outperforming the cases that used L1-W. This is expected since L1-W is unable to utilize a reference image. Case 1 and Case 2 had the best recovery of the phantom image out of all methods. Case 5, which also uses the same adaptive sampling density as cases 1 and 2, also shows accurate recovery. We notice that for cases with f_R (Case 1, Case 2, and Case 5) and f_{ND} (Case 17, Case 18, and Case 19) as their adaptive density sampling methods, the RSNR is high when using the LACS-MRI algorithm. f_{VD} combined with f_R performs better than f_{VDS} combined with f_R. Accordingly, we see that the adaptive density function plays a key role in the LACS-MRI algorithm.

Note that when we use only variable density sampling, we get the worst performance. Again, this is not surprising since this means we do not benefit from the use of a reference image. Recovery can be improved by using f_A but it does not compare favorably with f_R or f_{ND}. On the other hand, L1-W recovered the image poorly. Surprisingly, purely adaptive cases with f_R and f_{ND} (Case 13 and Case 14) performed the best among all combinations. In L1-W, we initialize f_{ND} to be the same as f_R which is why both show almost the exact same performance.

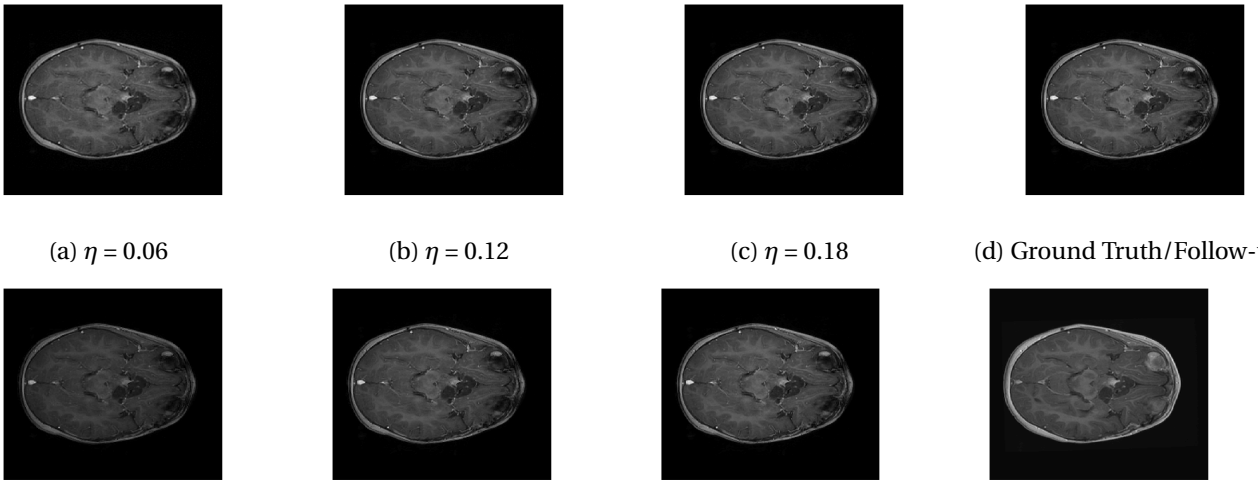

(a) $\eta = 0.06$ (b) $\eta = 0.12$ (c) $\eta = 0.18$ (d) Ground Truth/Follow-up

(e) $\eta = 0.06$ (f) $\eta = 0.12$ (g) $\eta = 0.18$ (h) Reference

Figure 2: Image reconstruction from Case 20 (a, b, and c) and Case 3 (e, f, and g) of a brain image. Case 20 is the best performing algorithm that uses L1-W and case 3 is the best performing algorithm that uses LACS-MRI

The phantom image that we used in the previous experiment is flat and has less texture while the brain images has a lot of wrinkles and complex shapes. While the phantom image is small, the brain image is too large to calculate f_{ND} efficiently. Thus, we conducted an experiment in the same setting as the phantom image excluding cases using f_A: cases 6 through 8 and 14 through 16. In this experiment, contrary to results for the phantom image experiment, L1-W performs better than LACS-MRI. For example, Case 3 performed the best among cases using LACS-MRI but it performs worse than 6 cases that use L1-W. Interestingly, one of the cases that shows an outstanding image recovery purely uses f_{VDS} with L1-W. However, purely using f_{VD} with L1-W shows a dramatically worse performance. In conclusion, we are unsure why each combination of mixed PDFs and reconstruction algorithm perform differently based on the images processed, but this differing behavior is crucial to highlight for practical applications. We conjecture that the texture of the image is one factor that decides the performance of each algorithm. In the future, we may explore this conjecture and analyze why this happens.

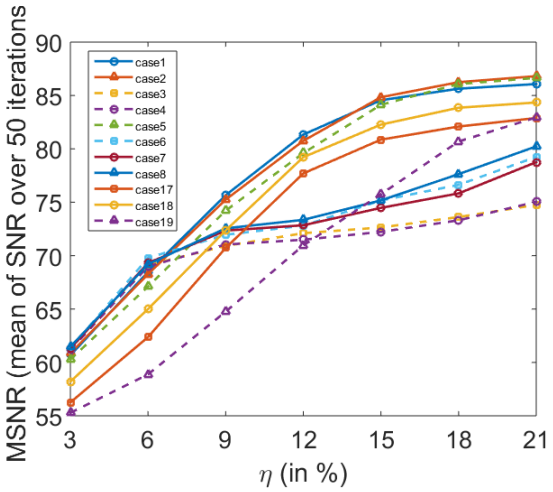

(a) Phantom image with high MSNR cases

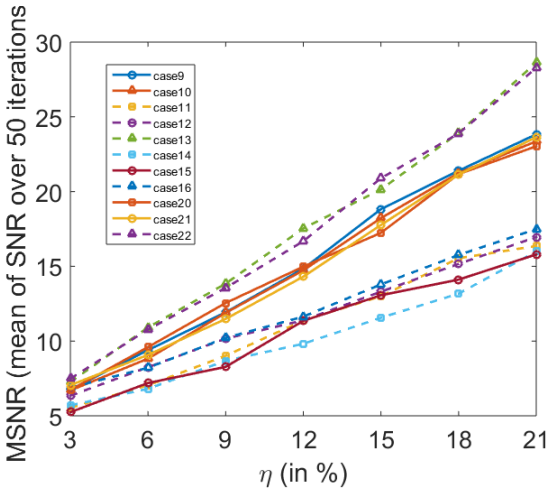

(b) Phantom image with low MSNR cases

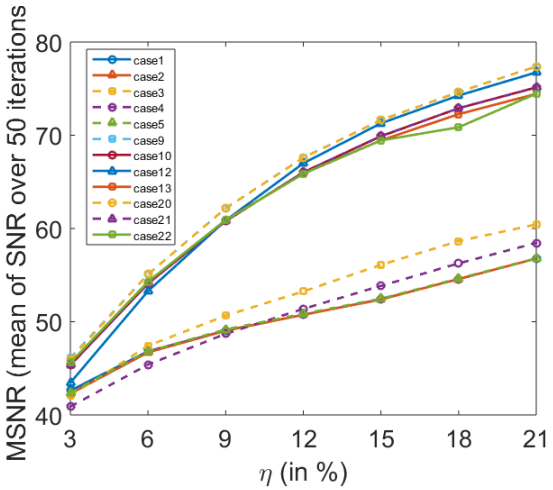

(c) Brain image with high MSNR cases

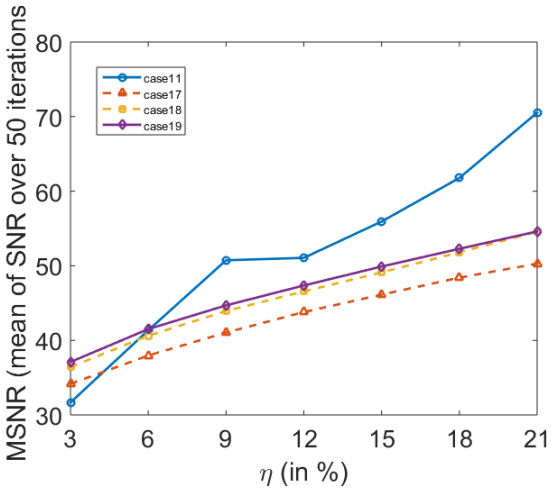

(d) Brain image with low MSNR cases

Figure 3: Summary of Phantom and Brain image experiments

3.2 Optimal parameters for sampling probability density functions

We next present simulations in which we tested different values of p and C in Krahmer and Ward's [12] variable density sampling function, which we will again refer to here as f_{VDS}. Recall f_{VDS} is defined via

$$\text{Prob}[(\omega_1^j, \omega_2^j) = (k_1, k_2)] = C' \min\left(C, \frac{1}{(k_1^2 + k_2^2)^p}\right),$$

where ω_i^j are frequencies, k_i are k-space coordinates, and C', C, and p are constants. In our first experiment, we measured RSNR for $.02 \leq p \leq 1.5$ at increments of .02, and $C = 1, .1, .01$, and .001, in an attempt to find optimal values of C; all experiments were performed on the brain image as above. We used the LACS-MRI algorithm to reconstruct the image and, in keeping with Krahmer and Ward's methods, set the compression level $\eta = .06$.

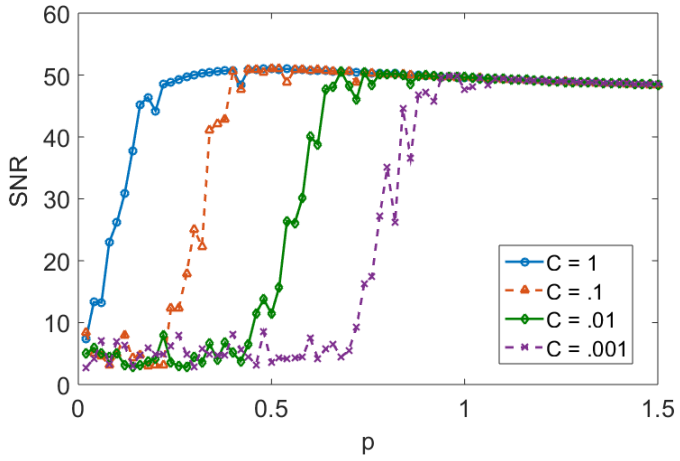

Figure 4: Reconstruction (by LACS-MRI) RSNR using f_{VDS} with different values of C and p.

The results are shown in Figure 4. We found that the larger C is, the greater the RSNR of the reconstruction is. However, as p becomes sufficiently large, there is virtually no difference in performance among the difference values of C. This is entirely consistent with how f_{VDS} is defined. An intuitive way to visualize f_{VDS} is to imagine the three-dimensional graph of $\frac{1}{(k_1^2+k_2^2)^p}$ where k_1 and k_2 are independent and p is fixed. The effect of selecting the smaller value of C and $\frac{1}{(k_1^2+k_2^2)^p}$ is to place an upper bound on $\frac{1}{(k_1^2+k_2^2)^p}$, which would otherwise blow up near the origin. f_{VDS} is precisely the two-dimensional intersection of $\frac{1}{(k_1^2+k_2^2)^p}$ and the plane $z = C$. In our earlier discussion about variable density sampling, we argue that sampling primarily near the origin in k-space yields a superior reconstruction. Conversely, a probability density function that samples uniformly at random everywhere will perform poorly. When p is small, $\frac{1}{(k_1^2+k_2^2)^p}$ is not very steep, even as $(k_1, k_2) \to 0$. Graphically, f_{VDS} has a large diameter and therefore will produce a lower RSNR image reconstruction. Furthermore, since $\frac{1}{(k_1^2+k_2^2)^p}$ is not steep, the difference in the probability density function for each value of C is drastic, which explains why higher values of C perform much better than lower values of C for low p. However, when p is large, $\frac{1}{(k_1^2+k_2^2)^p}$ becomes very steep near the origin, which results in both a small probability density function and a similar performance for each C.

Next, we tested the effects on RSNR by varying p while using mixed PDFs with f_{VDS}. We used different adaptive sampling functions and reconstruction algorithms. Specifically, we tested the LACS-MRI and ℓ_1-minimization with wavelet sparsity (L1-W) reconstruction algorithms with the adaptive-weighted reference based MRI (f_R), f_{VDS} only, f_A, and f_{VDS} and f_{ND} adaptive sampling functions. As discussed previously, LACS-MRI delivers an image reconstruction that is significantly superior to L1-W. We separated these eight cases into two graphs based on whether they were reconstructed using LACS-MRI or L_1-W, and thus whether they yielded high or low RSNR.

The results of this experiment are shown in Figure 5. In both the high RSNR and low RSNR cases, all combinations followed the same pattern as p varied, suggesting that the value of p is a significant factor in all reconstructions that utilize f_{VDS}. None of the low RSNR cases showed any discernible patterns as p varied. We conjecture that this is because the RSNR is simply too low to draw any meaningful results. On the other hand, reconstructions of the high RSNR cases do give meaningful results. First, although theoretical results of f_{VDS} suggest that the optimal value of p should be close to 1.5, we see here that instead the optimal value of p is around .7. However, also note that there is very little change in RSNR from the optimal p value and all other p values in a relatively large range. We conjecture that if p is within a sufficiently large range such as $.5 \leq p \leq 1.5$, the RSNR of the reconstruction will be bounded below, guaranteeing an accurate reconstruction. Examples of resconstructed images are shown in Figure 6, where there are only very minor discernable differences in reconstruction.

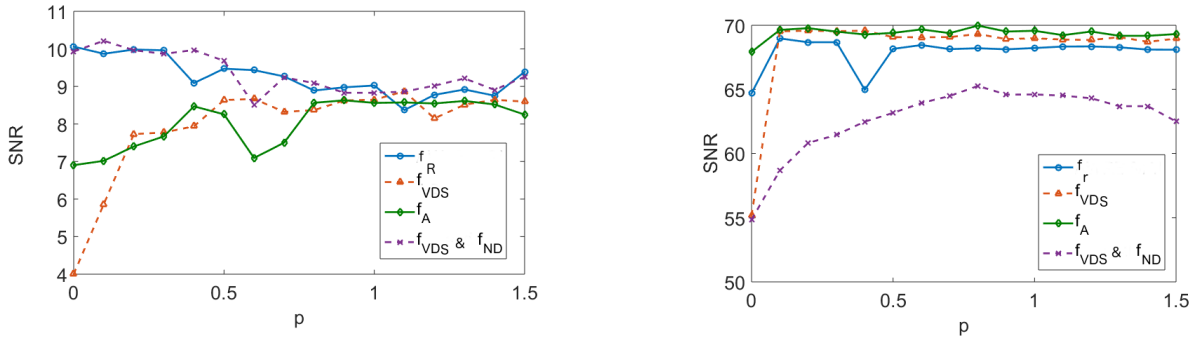

Figure 5: Left: Reconstruction SNR as a function of p using L1-W. Right: Reconstruction SNR as a function of p using LACS-MRI.

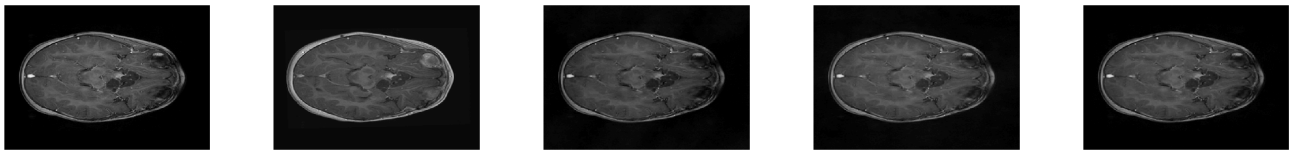

Figure 6: From left to right: the ground truth, reference image, and reconstructions for optimal p, $p = 1$ and $p = 1.5$.

3.3 Total Variation and gradient sparsification

Total variation (TV) minimization is a common process in image processing used for denoising and other applications. The *total variation* of an image $\boldsymbol{X}$ refers to the ℓ_1 norm of the discrete gradient, $\|\boldsymbol{X}\|_{TV} = \|\nabla \boldsymbol{X}\|_1$, where ∇ denotes the discrete gradient operator that concatenates the horizontal pixel-wise differences and the vertical differences. Suppose we have an NxN image, $\boldsymbol{X} \in \mathbb{C}^{NxN}$. Let $\boldsymbol{X}_{j,k}$ be the pixel in the j^{th} row and k^{th} column:

$$\boldsymbol{X} = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1N} \\ x_{21} & x_{22} & \dots & x_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ x_{N1} & x_{N2} & \dots & x_{NN} \end{bmatrix}. \tag{20}$$

We can define both the vertical and horizontal pixel-wise change by taking the respective directional derivatives:

$$(\boldsymbol{X}_x)_{j,k} = \begin{cases} \boldsymbol{X}_{j,k} - \boldsymbol{X}_{j+1,k} & 1 \le j \le N\,,\ 1 \le k < N-1 \\ \boldsymbol{X}_{j,k} & 1 \le j \le N\,,\ k = N \end{cases} \tag{21}$$

$$(\boldsymbol{X}_y)_{j,k} = \begin{cases} \boldsymbol{X}_{j,k} - \boldsymbol{X}_{j,k+1} & 1 \le j < N-1\,,\ 1 \le k \le N \\ \boldsymbol{X}_{j,k} & j = N\,,\ 1 \le k \le N \end{cases} \tag{22}$$

For example, we can write the vertical discrete gradient transform as a matrix in the form:

$$\boldsymbol{G} = \begin{bmatrix} 1 & -1 & 0 & \dots & \dots & 0 \\ 0 & 1 & -1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & 0 \\ \vdots & & & \ddots & \ddots & -1 \\ 0 & \dots & \dots & \dots & 0 & 1 \end{bmatrix}, \tag{23}$$

such that the vertical and horizontal gradients of $\boldsymbol{X}$ are $\nabla_y = \boldsymbol{GX}$ and $\nabla_x = \boldsymbol{GX}^T$. We then write $\nabla \boldsymbol{X} = (\nabla_x, \nabla_y)$.

When the total variation of an image is minimized, noise is removed and texture is smoothed out while the edges and important details are maintained. In the context of compressed sensing, the discrete gradient can be used as a sparsifying basis. It was proven by Needell and Ward [16] that even with undersampling, total variation minimization can recover an image robustly. Indeed, many natural images do not vary much in intensity from pixel to pixel which makes them compressible with respect to their discrete gradients. For example, the Shepp-Logan Phantom image is sparser with respect to its discrete gradient than its wavelet representation, see Figure 7.

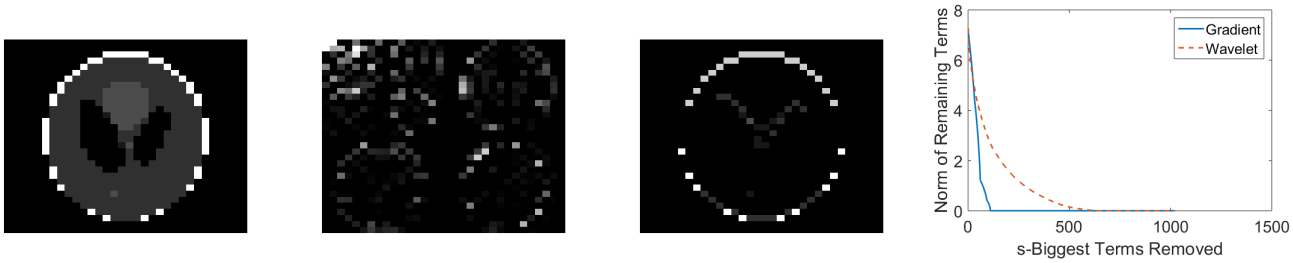

Figure 7: From left to right: original 32 x 32 Shepp-Logan Phantom image, wavelet representation, discrete gradient, and plot showing the decay of the Frobenius norm of the gradient/wavelet representation after removing large terms.

Due to its wide use in image processing, we consider the use of gradient sparsification with LACS-MRI. We show below that it is possible to produce image reconstructions with significantly higher Reconstruction SNR using gradient sparsification. We have repeated some of the simulations from above for different compression levels, η, using the gradient transform to sparsify instead of the wavelet transform. We will refer to this as "LACS-MRI with gradient sparsity." We compared the performance of each sparsification method at each compression level η using LACS-MRI and then compared this to the L1-W method. Again, we use a reference image and a follow-up image that we hope to recover as in Figure 8. We consider here only the phantom image, as it serves as a prime candidate for gradient sparsification, and is small enough that evaluation of all PDFs is computationally feasible.

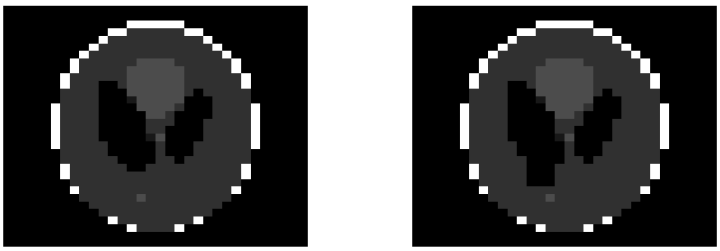

(a) (b)

Figure 8: (a) Reference phantom image and (b) Follow-up phantom image. Note the difference in the bottom of the left black spot; a "tumor" added by changing a few pixel intensities.

We re-ran simulations for cases 7, 8, 17, and 18 on the phantom. Over 30 trials, using the gradient on the phantom image resulted in much higher RSNR for every single case. Cases 7 and 17 both also outperformed cases 8 and 18 every time by about 13%. These cases both used the variable density sampling method f_{VD} and LACS-MRI. Figure 9 shows the reconstruction from Case 17. The recovered image using LACS-MRI with wavelet sparsity is more pixelated near the tumor while LACS-MRI with gradient sparsity produced an almost perfect reconstruction. Figure 11 shows the great disparity between the two sparsifying methods. On average, the four phantom reconstructions using the gradient only needed to sample about 6% of the data to achieve the same

RSNR that the reconstructions using the wavelets did at 21% sampling.

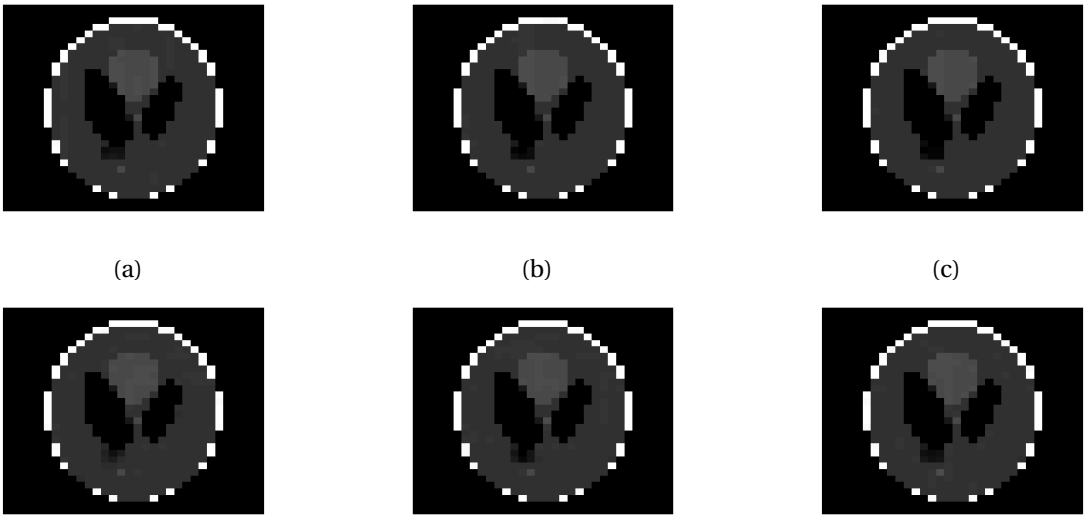

(a) (b) (c)

(d) (e) (f)

Figure 9: (Top) Case 17 using LACS-MRI with gradient sparsification. Reconstruction with compression levels (a) 6 (b) 12 and (c) 18. (Bottom) Case 17 using LACS-MRI with wavelet sparsification. Reconstruction with compression levels (d) 6 (b) 12 and (c) 18. The sampling methods used together were f_{VD} and f_{ND}.

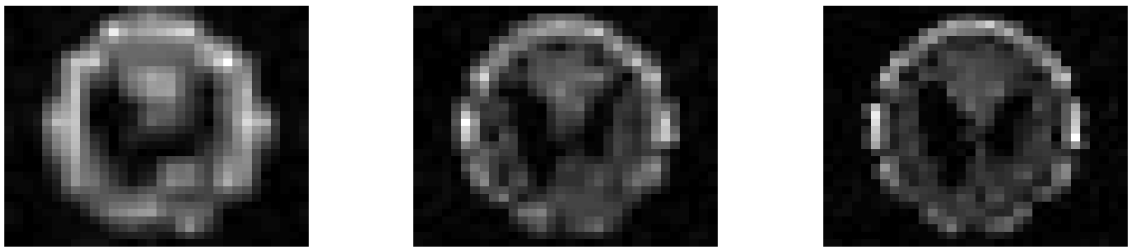

(a) (b) (c)

Figure 10: Case 10 using ℓ_1-minimization with wavelet sparsification. Reconstruction with compression levels (a) 6 (b) 12 and (c) 18. The sampling methods used together were f_{VDS} and f_R.

We also compare the LACS-MRI gradient method to the Cases 9 and 10 which use ℓ_1-minimization with wavelet sparsity. Clearly, the L1-W method recovered the phantom image very poorly, as shown in Figure 10 and Figure 11 (b). The image is blurry even when 18% of the available data was sampled. The RSNR of L1-W is worse than both the RSNR achieved by LACS-MRI with wavelet sparsity and LACS-MRI with gradient sparsity.

We can conclude from these experiments that different sparsification methods can clearly have a drastic effect on the reconstruction of the follow-up image. LACS-MRI with gradient sparsity is better for images with a sparser gradient than wavelet transform, such as the phantom image. For images with high amounts of texture we expect the wavelet transform to work better. Therefore, if we know that an image will be more sparse in one domain over the other, we can appropriately choose the reconstruction method. Such a determination of course may be non-trivial in some applications, and much research in compressed sensing and sparse approximation focuses on this judgement.

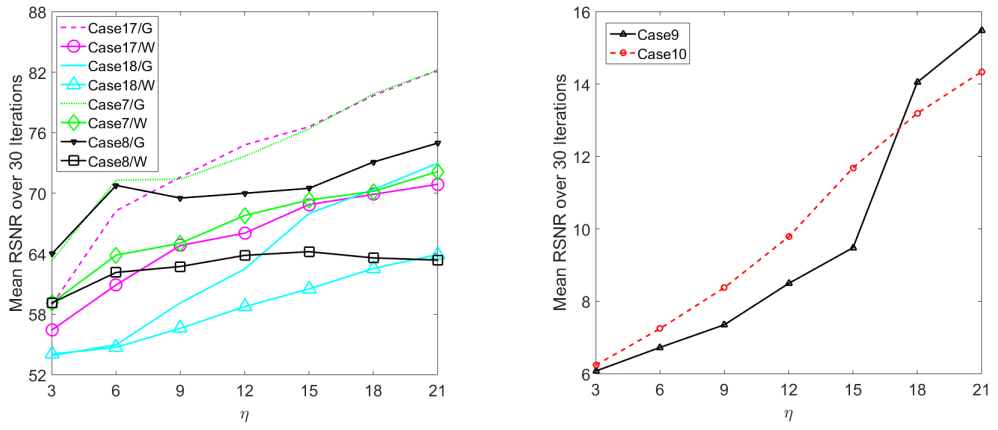

(a) (b)

Figure 11: (a) Reconstruction SNR of the recovered image at different compression levels for the phantom image ('G' and 'W' indicate the gradient and wavelet sparsifying transform used). (b) Reconstruction SNR of the recovered phantom image at different compression levels for cases 9 and 10 using L1-W (lower SNR). η is in %.

4 Grayscale Correction

In a clinical setting, longitudinal MRI images may sometimes be taken on different MRI machines, on different days, with possibly different contrast settings, e.g. (T1 weighted, T2 weighted, Proton Density weighted) [10]. This causes many practical issues in utilizing prior MRI scans as reference images; we tackle one of these problems, which is differing grayscale intensities. We consider the problem where intensities of pixels of the same tissue are scaled very differently, and thus have differing amounts of contrast. This scaling factor can result in making the difference term, Term 2 in Equation 9, be very large and thus the solution of the optimization problem would prioritize minimizing the difference term at the expense of fidelity to the actual measurements, resulting in low signal-to-noise ratio for the resultant image.

The effect of the scale difference would be most pronounced in the argument of Term 2, but also impacts $\boldsymbol{W}_1$ and $\boldsymbol{W}_2$, since the values of these weighting terms also depend on $\boldsymbol{x}-\boldsymbol{x}_0$. As shown in Figures 13 and 14, the intensity shift causes significantly reduced RSNR. We have designed an algorithm and a modification of reference based MRI to compensate for the shift in grayscale intensity and ensure high RSNR.

We first attempt a toy problem aimed at solving this issue, where we model the shift in grayscale intensity as linear, i.e. for some $c>0$, $\boldsymbol{X}=c\boldsymbol{X}_0$. If we know c apriori, we can solve the optimization problem (9) by modifying Term 2 in order to compensate for the intensity shift. Accordingly, Term 2 becomes $\lambda_2\|W_2(\boldsymbol{X}-c\boldsymbol{X}_0)\|_1$. This change corrects the inaccuracy of term 2 and allows the reconstruction to proceed based on the actual structure of the sampled image. We are now considering the correctly scaled difference between the reference and follow-up. The issue is how to estimate c.

As $\boldsymbol{X}=c\boldsymbol{X}_0$, we should have $c=\frac{\boldsymbol{X}}{\boldsymbol{X}_0}$. Since we only have access to the values of the Fourier coefficients obtained from the MRI, our proposed modified algorithm iteratively provides an estimate of c, based on the points sampled in each sampling round of LACS-MRI. The estimate of c is then used to evaluate $c\boldsymbol{X}_0$ in each iteration of the numerical method (rather than just $\boldsymbol{X}_0$) to solve the original optimization problem (9). We define $\boldsymbol{Y}=\boldsymbol{F}_u\boldsymbol{X}$ and $\boldsymbol{Y}_0=\boldsymbol{F}_u\boldsymbol{X}_0$ to be the Fourier transforms of the follow-up and the reference image respectively, and describe our algorithmic modification by the following pseudocode.

Algorithm 2 Grayscale Correction for Reference Based MRI (GSC)

Input: $\boldsymbol{Y}_{\text{samp}}$, the sampling locations of the follow-up image, $(\boldsymbol{Y}_0)_{\text{samp}}$, the sampling locations of the reference image c , the current grayscale estimate, i, number of sampling iterations

$c^{'} = \frac{\sum_{i,j} \boldsymbol{Y}_{i,j}}{\sum_{i,j} (\boldsymbol{Y}_0) i,j}$

Output: $c = c^{'} \frac{1}{i} + c \frac{i-1}{i}$

The following plots clearly show the effectiveness of this scale correction method and the noisy results when scale is left uncorrected. Scale Correction (SC) refers to the reconstructions using the GSC algorithm, while No Scale Correction (NSC) is the standard unmodified LACS-MRI.

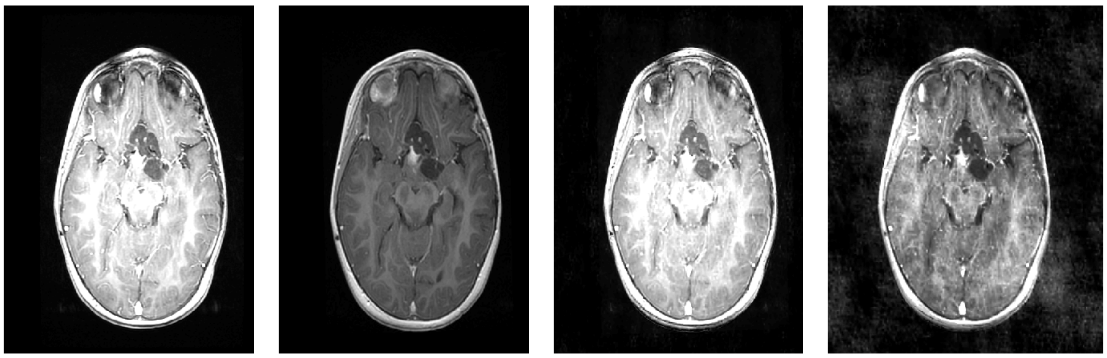

(a) Ground Truth (b) Reference (c) Scale Correction (d) No Scale Correction

Figure 12: (a) is the ground truth of the follow-up image, (b) is the reference image (with scaling parameter $c = 2$), (c) and (d) are reconstructions with η=0.06 over 3 sampling iterations using LACS-MRI.

Figure 12 shows the effect of a 2-fold scale difference on LACS-MRI image reconstructions. There are extensive streaking artifacts and significant detail is lost compared to the ground truth and the reconstruction using the modified method. Clearly, the Grayscale Compensation for Reference Based MRI Algorithm (GSC) can reduce these artifacts and help improve the visibility of structure.

Figure 13 shows the relationship of RSNR and c. Figure 13 is validated by the fact that when $c = 1$, the RSNR of both the Brain image and Shepp Logan Phantom converge to the same value. There is some residual error due to noise in the estimation of c.

Figure 14 shows the relationship between RSNR and η, where η is the total amount of the source image sampled. As to be expected, in both cases RSNR increases with increasing η, or more samples will yield better reconstructions.

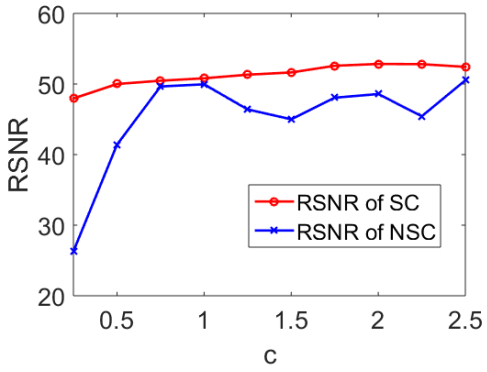

(a) Brain Image

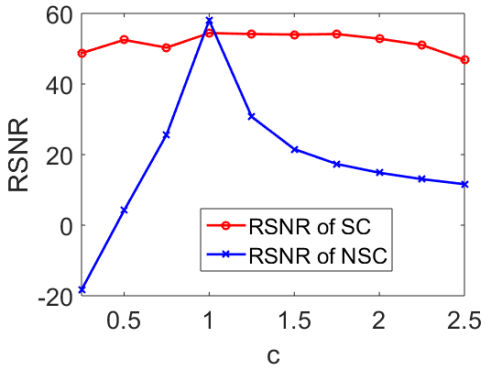

(b) Shepp Logan Phantom

Figure 13: Reconstruction SNR vs c for c= 0.25, 0.5, ... 2.5 with η=0.15 over 3 sampling iterations, each data point is the average of 50 trials.

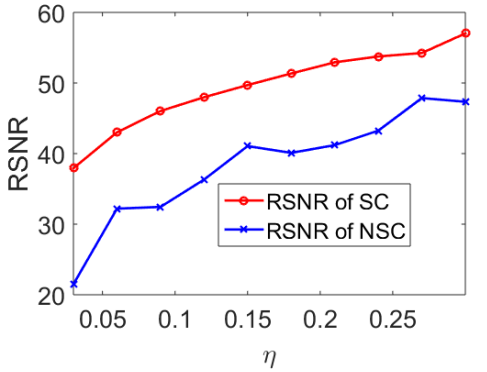

(a) Brain Image

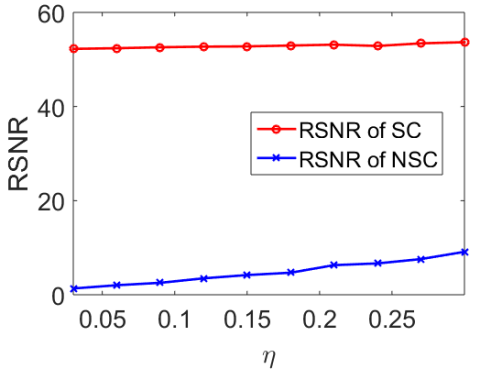

(b) Shepp Logan Phantom

Figure 14: Reconstruction SNR vs η for η= 0.03, 0.06, ... 0.30 over 3 sampling iterations with c=0.5, each data point is the average of 50 trials.

These results support the validity of our solution to this toy problem and let us hope that similar but less naive techniques can be applied to more complex issues, i.e. variable contrast and scaling localized to sub-regions of brain scans. As in the linear toy problem, subregions exhibiting this type of behavior can result in lowered RSNR. The solution to this larger problem would greatly enhance the value of LACS-MRI in clinical situations and for greater robustness to variation in MRI acquisition parameters.

5 Conclusion

We have built upon previous work done by Weizman et. al. [20] by changing the sampling method, applying gradient as a sparsifier, and compensating for differences in grayscale level between reference and follow-up images. In addition, we investigated the combinations of different adaptive sampling functions, variable sampling functions, and reconstruction algorithms in an attempt to find an optimal image reconstruction method. We found that the accuracy of resultant reconstructions is dependent on the properties of the image itself, in a certain basis, and that taking advantage of this knowledge can yield high RSNR given minimal samples.

We hope to pursue future research in more complex non-linear grayscale compensation issues where only a portion of the image is scaled differently. We also hope to find theoretical bounds on these reference based methods applied to different classes of images.

References

[1] BERMAN, K., HALPERN, H., KAFTAL, V., AND WEISS, G. Matrix norm inequalities and the relative dixmier property. *Integral Equations and Operator Theory 11*, 1 (1988), 28–48.

[2] CANDES, E., RUDELSON, M., TAO, T., AND VERSHYNIN, R. Error correction via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)* (2005), IEEE, pp. 668–681.

[3] CANDÈS, E. J. Compressive sampling. In *Proc. Int. Congress of Mathematics* (Madrid, Spain, 2006), vol. 3, pp. 1433–1452.

[4] CANDÈS, E. J., ROMBERG, J., AND TAO, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory 52*, 2 (2006), 489–509.

[5] CANDES, E. J., ROMBERG, J. K., AND TAO, T. Stable signal recovery from incomplete and inaccurate measurements. *Communications on pure and applied mathematics 59*, 8 (2006), 1207–1223.

[6] DAUBECHIES, I. *Ten lectures on wavelets*. SIAM, Philadelphia, PA, 1992.

[7] DAVENPORT, M. A., MASSIMINO, A. K., NEEDELL, D., AND WOOLF, T. Constrained adaptive sensing. *arXiv preprint arXiv:1506.05889* (2015).

[8] DONOHO, D. L. Compressed sensing. *IEEE T. Inform. Theory 52*, 4 (Apr. 2006), 1289–1306.

[9] HOLLINGWORTH, W., TODD, C. J., BELL, M. I., ARAFAT, Q., GIRLING, S., KARIA, K. R., AND DIXON, A. K. The diagnostic and therapeutic impact of mri: an observational multi-centre study. *Clinical radiology 55*, 11 (2000), 825–831.

[10] HUANG, J., CHEN, C., AND AXEL, L. Fast multi-contrast mri reconstruction. *Magnetic resonance imaging 32*, 10 (2014), 1344–1352.

[11] KIRSCHEN, D. G. Use in environment of an mri scanner, Feb. 14 1989. US Patent 4,804,261.

[12] KRAHMER, F., AND WARD, R. Stable and robust sampling strategies for compressive imaging. *Image Processing, IEEE Transactions on 23*, 2 (2014), 612–622.

[13] LUSTIG, M., DONOHO, D., AND PAULY, J. M. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic resonance in medicine 58*, 6 (2007), 1182–1195.

[14] LUSTIG, M., AND PAULY, J. M. Spirit: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magnetic Resonance in Medicine 64*, 2 (2010), 457–471.

[15] MUTHUKRISHNAN, S. *Data streams: Algorithms and applications*. Now Publishers Inc, 2005.

[16] NEEDELL, D., AND WARD, R. Stable image reconstruction using total variation minimization. *SIAM Journal on Imaging Sciences 6*, 2 (2013), 1035–1058.

[17] RAUHUT, H., SCHNASS, K., AND VANDERGHEYNST, P. Compressed sensing and redundant dictionaries. *IEEE Trans. Inform. Theory 54*, 5 (2008), 2210–2219.

[18] RUDELSON, M., AND VERSHYNIN, R. On sparse reconstruction from fourier and gaussian measurements. *Comm. Pure Appl. Math. 61* (2008), 1025–1045.

[19] WEIZMAN, L., ELDAR, Y. C., AND BASHAT, D. B. The application of compressed sensing for longitudinal mri. *arXiv preprint arXiv:1407.2602* (2014).

[20] WEIZMAN, L., ELDAR, Y. C., AND BASHAT, D. B. Compressed sensing for longitudinal mri: An adaptive-weighted approach. *Medical physics 42*, 9 (2015), 5195–5208.